

1 **CLAIMS**

2
3 1. One or more processor-accessible media comprising processor-
4 executable instructions that, when executed, direct a device to perform actions
5 comprising:

6 accepting a connection;

7 aggregating a connection state for the connection from a protocol stack; and

8 sending the connection state.
9

10 2. The one or more processor-accessible media as recited in claim 1,
11 wherein the action of accepting comprises an action of:

12 sending an acknowledgment packet in response to a connection-
13 requesting packet.
14

15 3. The one or more processor-accessible media as recited in claim 1,
16 comprising the processor-executable instructions that, when executed, direct the
17 device to perform a further action comprising:

18 receiving data for the connection;

19 wherein the action of aggregating comprises an action of:

20 aggregating the connection state from a protocol state of the protocol
21 stack and the data.
22
23
24
25

1 4. The one or more processor-accessible media as recited in claim 1,
2 wherein the action of aggregating comprises an action of:

3 compiling a protocol state from the protocol stack.
4

5 5. The one or more processor-accessible media as recited in claim 4,
6 wherein the action of compiling comprises an action of:

7 compiling the protocol state from the protocol stack starting
8 at a highest level of the protocol stack.
9

10 6. The one or more processor-accessible media as recited in claim 4,
11 wherein the action of compiling comprises an action of:

12 compiling the protocol state from the protocol stack at a
13 transmission control protocol (TCP) stack portion and an internet
14 protocol (IP) stack portion.
15

16 7. The one or more processor-accessible media as recited in claim 1,
17 wherein the action of sending comprises actions of:

18 bundling the connection state with a flow identifier that corresponds
19 to the connection to produce a binary blob; and

20 transmitting the binary blob from an originating device to a target
21 device.
22
23
24
25

1 8. The one or more processor-accessible media as recited in claim 1,
2 wherein the action of sending comprises actions of:

3 bundling the connection state with a flow identifier that corresponds
4 to the connection to produce a binary blob; and

5 transmitting the binary blob from an originating device to a target
6 device in a reliable manner such that the binary blob may be received intact
7 at the target device even if one or more packets that comprise the binary
8 blob are lost or corrupted.

9
10 9. The one or more processor-accessible media as recited in claim 1,
11 comprising the processor-executable instructions that, when executed, direct the
12 device to perform further actions comprising:

13 selecting a flow identifier for the connection responsive to a connection
14 counter; and

15 sending the flow identifier to identify packets corresponding to the
16 connection.

17
18 10. The one or more processor-accessible media as recited in claim 1,
19 wherein the action of sending comprises an action of:

20 sending the connection state to a targeted device;

21 wherein the processor-executable instructions, when executed, direct the
22 device to perform a further action comprising:

23 forwarding subsequent packets for the connection to the targeted device
24 using a flow identifier to encapsulate the subsequent packets.

25

1 **11.** One or more processor-accessible media comprising processor-
2 executable instructions that, when executed, direct a device to perform actions
3 comprising:

4 receiving a connection state for a connection;
5 injecting the connection state for the connection into a network stack; and
6 continuing the connection using the injected connection state.

7
8 **12.** The one or more processor-accessible media as recited in claim 11,
9 wherein the action of continuing comprises an action of:

10 continuing the connection by indicating received packets up to an
11 application in accordance with the injected connection state.

12
13 **13.** The one or more processor-accessible media as recited in claim 11,
14 wherein:

15 the action of receiving comprises an action of:

16 receiving the connection state, the connection state having a protocol
17 state and data for the connection; and

18 the action of injecting comprises an action of:

19 injecting the protocol state into a protocol stack portion of the
20 network stack.

1 **14.** The one or more processor-accessible media as recited in claim 13,
2 wherein the action of injecting the connection state further comprises an action of:
3 indicating the data for the connection up the network stack toward an
4 application.

5
6 **15.** The one or more processor-accessible media as recited in claim 11,
7 wherein the action of injecting comprises an action of:
8 infusing a protocol state from the connection state into a protocol
9 stack portion of the network stack.

10
11 **16.** The one or more processor-accessible media as recited in claim 15,
12 wherein the action of infusing comprises an action of:
13 infusing the protocol state into the protocol stack starting at a
14 highest level of the protocol stack.

15
16 **17.** The one or more processor-accessible media as recited in claim 15,
17 wherein the action of infusing comprises an action of:
18 infusing the protocol state into the protocol stack at a
19 transmission control protocol (TCP) stack portion and an internet
20 protocol (IP) stack portion.

1 **18.** The one or more processor-accessible media as recited in claim 11,
2 wherein the action of receiving comprises actions of:

3 receiving a binary blob from an originating device at a target device,
4 the binary blob including the connection state and a flow identifier that
5 corresponds to the connection; and

6 unbundling the connection state and the flow identifier at a level of
7 the network stack that is below a protocol stack portion of the network
8 stack.

9
10 **19.** The one or more processor-accessible media as recited in claim 11,
11 comprising the processor-executable instructions that, when executed, direct the
12 device to perform further actions comprising:

13 receiving an encapsulation mapping; and

14 storing the received encapsulation mapping in an encapsulation mapping
15 table that may be accessed according to flow identifier.

1 **20.** The one or more processor-accessible media as recited in claim 11,
2 wherein the action of receiving comprises an action of:

3 receiving the connection state from an originating device;

4 wherein the processor-executable instructions, when executed, direct the
5 device to perform a further action comprising:

6 receiving from the originating device encapsulated packets that have a flow
7 identifier; and

8 de-encapsulating the encapsulated packets using an encapsulation mapping
9 entry that links the flow identifier to a source/destination pair.

10
11 **21.** One or more processor-accessible media comprising processor-
12 executable instructions that comprise at least part of an application, wherein the
13 processor-executable instructions, when executed, enable the application to initiate
14 a connection migration by issuing a migrate connection function call to a socket
15 layer such that the socket layer propagates the migrate connection function call to
16 a protocol stack, the migrate connection function call precipitating at the protocol
17 stack a compilation of protocol state from the protocol stack.

18
19 **22.** The one or more processor-accessible media as recited in claim 21,
20 wherein the application comprises at least one of a general application, a hosted
21 application, and a load balancing application.

1 **23.** A device comprising:

2 a connection migrator that is configured to migrate connections away from
3 the device; the connection migrator capable of precipitating a compilation of
4 protocol state for a connection across a protocol stack; the connection migrator
5 adapted to aggregate the compiled protocol state with data for the connection into
6 an aggregated connection state of the connection; the connection migrator further
7 capable of causing the aggregated connection state to be sent toward a target
8 device.

9
10 **24.** The device as recited in claim 23, wherein the connection migrator
11 is realized at least partially in software.

12
13 **25.** The device as recited in claim 23, wherein the device further
14 comprises:

15 a classifier that is capable of accepting the connection, the classifier
16 adapted to issue a migrate connection command to the connection migrator.

17
18 **26.** The device as recited in claim 25, wherein the data, which is
19 aggregated with the compiled protocol state into the aggregated connection state,
20 comprises data that has been acknowledged by the classifier.

1 **27.** The device as recited in claim 23, wherein the connection migrator
2 comprises:

3 a migrator shim that is located above the protocol stack in a network
4 stack of the device; and

5 a migrator intermediate driver that is located below the protocol
6 stack in the network stack of the device.

7
8 **28.** The device as recited in claim 27, wherein:

9 the migrator shim is located between the protocol stack and a socket
10 layer of the network stack of the device; and

11 the migrator intermediate driver is located between the protocol
12 stack and at least one miniport of the network stack of the device.

13
14 **29.** The device as recited in claim 28, wherein the migrator intermediate
15 driver is located at a protocol-hardware interface layer of the network stack of the
16 device.

17
18 **30.** The device as recited in claim 27, wherein:

19 the migrator shim is adapted to receive a migrate connection
20 command and to propagate the migrate connection command to the
21 protocol stack; and

22 the migrator intermediate driver is adapted to divert a copy of the
23 data for the connection for subsequent aggregation with the compiled
24 protocol state.

1 **31.** The device as recited in claim 27, wherein the migrator intermediate
2 driver is adapted to bundle the aggregated connection state and a flow identifier
3 for the connection into a binary blob.

4
5 **32.** The device as recited in claim 23, wherein the connection comprises
6 a transmission control protocol/internet protocol (TCP/IP) connection, and the
7 compiled protocol state includes information related to a TCP/IP connection.

8
9 **33.** The device as recited in claim 23, wherein the device further
10 comprises:

11 the protocol stack, the protocol stack adapted to compile the protocol state
12 for the connection into the compiled protocol state responsive to a migrate
13 connection command.

14
15 **34.** A device comprising:

16 a connection migrator that is configured to migrate connections onto the
17 device; the connection migrator adapted to intercept a connection state for a
18 connection that is sent from an originating device, the connection state including a
19 protocol state and data; the connection migrator further adapted to inject the
20 connection state into a network stack of the device; the connection migrator
21 capable of precipitating an infusion of the protocol state for the connection across
22 a protocol stack of the network stack.

1 **35.** The device as recited in claim 34, wherein the connection migrator
2 comprises at least part of load balancing infrastructure that is resident at and
3 executing on the device.

4
5 **36.** The device as recited in claim 34, wherein the device further
6 comprises:

7 an application that is hosted on the device;

8 wherein the connection migrator is further configured to migrate
9 connections onto the device in a manner that is transparent to the application.

10
11 **37.** The device as recited in claim 34, wherein the data included in the
12 connection state comprises data that was acknowledged at the originating device.

13
14 **38.** The device as recited in claim 34, wherein the connection migrator
15 comprises:

16 a migrator shim that is located above the protocol stack in the
17 network stack of the device; and

18 a migrator intermediate driver that is located below the protocol
19 stack in the network stack of the device.

1 **39.** The device as recited in claim 38, wherein:

2 the migrator shim is located between the protocol stack and a socket
3 layer of the network stack of the device; and

4 the migrator intermediate driver is located between the protocol
5 stack and at least one miniport of the network stack of the device.

6
7 **40.** The device as recited in claim 38, wherein:

8 the migrator shim is adapted to be notified by the migrator
9 intermediate driver of a migration uploading procedure and is adapted to
10 initiate a protocol state infusion routine with the protocol stack; and

11 the migrator intermediate driver is adapted to detect arrival of the
12 connection state for the connection, to divert the connection state away
13 from a lower portion of the protocol stack, and to notify the migrator shim
14 of the migration uploading procedure.

15
16 **41.** The device as recited in claim 38, wherein the migrator intermediate
17 driver is adapted to un-bundle a received binary blob that includes the connection
18 state for the connection and a flow identifier corresponding to the connection.

19
20 **42.** The device as recited in claim 34, wherein the device further
21 comprises:

22 the protocol stack, the protocol stack including a transmission control
23 protocol (TCP) layer and an internet protocol (IP) layer.

1 **43.** The device as recited in claim 34, wherein the device further
2 comprises:

3 the protocol stack, the protocol stack adapted to infuse the protocol state for
4 the connection across the protocol stack responsive to initiation of an infuse
5 protocol state routine.

6
7 **44.** An arrangement for connection manipulation, comprising:
8 migration means for migrating connections from an originating device to a
9 target device by transferring a connection state for a connection being migrated
10 from the originating device to the target device; and

11 tunneling means for tunneling packets for migrated connections in an
12 encapsulated format from the originating device to the target device.

13
14 **45.** The arrangement as recited in claim 44, wherein the migration
15 means comprises:

16 aggregation means for aggregating the connection state from a
17 protocol state and from data acknowledged for the connection being
18 migrated.

19
20 **46.** The arrangement as recited in claim 45, wherein the aggregation
21 means comprises:

22 compilation means for compiling the protocol state from a
23 protocol stack.

1 **47.** The arrangement as recited in claim 44, wherein the tunneling
2 means comprises:

3 encapsulation means for encapsulating the packets for the migrated
4 connections using flow identifiers that replace at least a part of
5 source/destination address information pairs of the packets.
6

7 **48.** The arrangement as recited in claim 44, wherein the tunneling
8 means comprises:

9 mapping means for mapping individual source/destination address
10 information pairs to individual flow identifiers, respective
11 source/destination address information pairs identifying respective migrated
12 connections.
13

14 **49.** The arrangement as recited in claim 44, wherein the migration
15 means comprises:

16 driver means for bundling the connection state with a flow identifier
17 that is usable for encapsulating packets of the connection being migrated.
18

19 **50.** The arrangement as recited in claim 44, wherein the migration
20 means comprises:

21 injection means for injecting a protocol state portion of the
22 connection state for the connection being migrated into a protocol stack
23 portion of a network stack.
24
25

1 **51.** The arrangement as recited in claim 44, wherein the tunneling
2 means comprises:

3 de-encapsulation means for de-encapsulating the tunneled packets
4 for the migrated connections using flow identifiers that link to at least a part
5 of source/destination address information pairs of the packets.
6

7 **52.** The arrangement as recited in claim 44, wherein the tunneling
8 means comprises:

9 mapping means for mapping individual flow identifiers to individual
10 source/destination address information pairs, respective source/destination
11 address information pairs identifying respective migrated connections.
12

13 **53.** The arrangement as recited in claim 44, wherein the migration
14 means comprises:

15 driver means for diverting the tunneled packets for the migrated
16 connections in the encapsulated format and directing them to the tunneling
17 means.
18

19 **54.** The arrangement as recited in claim 44, wherein the arrangement
20 comprises at least one device.
21

22 **55.** The arrangement as recited in claim 44, wherein the arrangement
23 comprises one or more processor-accessible media.
24
25

1 **56.** A device comprising:

2 a tunneler that is configured to tunnel packets away from the device; the
3 tunneler having access to an encapsulation mapping table, the encapsulation
4 mapping table including a plurality of encapsulation mapping entries, each
5 encapsulation mapping entry linking at least a portion of a source/destination pair
6 to a flow identifier; the tunneler adapted to accept a packet having a particular
7 source/destination pair; the tunneler capable of looking up a particular flow
8 identifier at a particular encapsulation mapping entry using at least a portion of the
9 particular source/destination pair; wherein the tunneler is further adapted to
10 encapsulate the packet by replacing part of the packet with the particular flow
11 identifier.

12
13 **57.** The device as recited in claim 56, wherein the tunneler is realized at
14 least partially in software.

15
16 **58.** The device as recited in claim 56, wherein the device further
17 comprises:

18 a forwarder that is capable of receiving the packet from a client, the
19 forwarder adapted to provide the packet to the tunneler.

20
21 **59.** The device as recited in claim 56, wherein the tunneler is further
22 adapted to encapsulate the packet by replacing at least part of the particular
23 source/destination pair of the packet with the particular flow identifier.

1 **60.** The device as recited in claim 56, wherein the flow identifier
2 comprises 32 bits.

3
4 **61.** The device as recited in claim 56, wherein the particular
5 source/destination pair comprises a particular source/destination address
6 information pair.

7
8 **62.** A device comprising:
9 a tunneler that is configured to tunnel packets into the device; the tunneler
10 having access to an encapsulation mapping table, the encapsulation mapping table
11 including a plurality of encapsulation mapping entries, each encapsulation
12 mapping entry linking a flow identifier to at least a portion of a source/destination
13 pair; the tunneler adapted to accept an encapsulated packet having a particular
14 flow identifier; the tunneler capable of looking up a particular source/destination
15 pair at a particular encapsulation mapping entry using the particular flow
16 identifier; wherein the tunneler is further adapted to de-encapsulate the
17 encapsulated packet by replacing the particular flow identifier with at least part of
18 the particular source/destination pair.

19
20 **63.** The device as recited in claim 62, wherein the tunneler comprises a
21 migrator intermediate driver that is adapted to intercept the encapsulated packet to
22 prevent the encapsulated packet from being provided to a protocol stack of the
23 device.

1 **64.** The device as recited in claim 62, wherein the device further
2 comprises:

3 a protocol stack;

4 wherein the tunneler is further adapted to de-encapsulate the encapsulated
5 packet by replacing the particular flow identifier with at least part of the particular
6 source/destination pair to produce a de-encapsulated packet; the tunneler capable
7 of indicating up the de-encapsulated packet to the protocol stack.

8
9 **65.** The device as recited in claim 64, wherein the tunneler comprises a
10 virtual network adapter that is bound to a destination address of the de-
11 encapsulated packet.

12
13 **66.** The device as recited in claim 62, wherein the device comprises a
14 host of a cluster of hosts.

15
16 **67.** The device as recited in claim 62, wherein the particular
17 source/destination pair corresponds to a transmission control protocol/internet
18 protocol (TCP/IP) 4-tuple; and wherein the tunneler is further adapted to de-
19 encapsulate the encapsulated packet by replacing the particular flow identifier with
20 a source TCP port and a destination TCP port from the TCP/IP 4-tuple.

1 **68.** One or more processor-accessible media comprising processor-
2 executable instructions that, when executed, direct a device to perform actions
3 comprising:

4 obtaining at least a portion of a source/destination pair from an incoming
5 packet;

6 accessing an encapsulation mapping table using the obtained at least a
7 portion of the source/destination pair to locate an encapsulation mapping entry;

8 extracting a flow identifier from the located encapsulation mapping entry;

9 and

10 replacing part of the incoming packet with the extracted flow identifier to
11 produce an encapsulated packet.

12
13 **69.** The one or more processor-accessible media as recited in claim 68,
14 comprising the processor-executable instructions that, when executed, direct the
15 device to perform a further action comprising:

16 receiving the incoming packet from a client.

17
18 **70.** The one or more processor-accessible media as recited in claim 68,
19 comprising the processor-executable instructions that, when executed, direct the
20 device to perform a further action comprising:

21 routing the encapsulated packet to a host.

1 71. The one or more processor-accessible media as recited in claim 68,
2 wherein the source/destination pair of the incoming packet comprises a
3 source/destination address information pair of the incoming packet.

4
5 72. The one or more processor-accessible media as recited in claim 68,
6 wherein at least a portion of the processor-executable instructions comprise at
7 least part of load balancing infrastructure.

8
9 73. The one or more processor-accessible media as recited in claim 68,
10 wherein the action of replacing comprises an action of:

11 replacing at least part of the source/destination pair with the
12 extracted flow identifier to produce the encapsulated packet.

13
14 74. The one or more processor-accessible media as recited in claim 68,
15 wherein:

16 the action of obtaining comprises an action of:

17 obtaining at least a portion of a transmission control
18 protocol/internet protocol (TCP/IP) 4-tuple from the incoming packet; and
19 the action of replacing comprises an action of:

20 replacing a source TCP port and a destination TCP port from the
21 incoming packet with the extracted flow identifier.

1 75. The one or more processor-accessible media as recited in claim 74,
2 comprising the processor-executable instructions that, when executed, direct the
3 device to perform a further action comprising:

4 replacing a source IP address and a destination IP address from the
5 incoming packet with an IP address of an originating device and an IP address of a
6 target device, respectively.

7
8 76. One or more processor-accessible media comprising processor-
9 executable instructions that, when executed, direct a device to perform actions
10 comprising:

11 obtaining a flow identifier from an encapsulated packet;

12 accessing an encapsulation mapping table using the obtained flow identifier
13 to locate an encapsulation mapping entry;

14 extracting at least a portion of a source/destination pair from the located
15 encapsulation mapping entry, the source/destination pair identifying a flow of
16 packets for a connection; and

17 replacing the flow identifier of the encapsulated packet with the extracted at
18 least a portion of the source/destination pair to produce a de-encapsulated packet
19 for the connection.

20
21 77. The one or more processor-accessible media as recited in claim 76,
22 comprising the processor-executable instructions that, when executed, direct the
23 device to perform a further action comprising:

24 receiving the encapsulated packet from load balancing infrastructure.
25

1 **78.** The one or more processor-accessible media as recited in claim 76,
2 comprising the processor-executable instructions that, when executed, direct the
3 device to perform a further action comprising:

4 indicating the de-encapsulated packet up to a protocol stack.
5

6 **79.** The one or more processor-accessible media as recited in claim 78,
7 wherein the action of indicating comprises an action of:

8 indicating the de-encapsulated packet up on a virtual network
9 adapter that is bound to a destination address of the de-encapsulated packet.
10

11 **80.** The one or more processor-accessible media as recited in claim 76,
12 comprising the processor-executable instructions that, when executed, direct the
13 device to perform further actions comprising:

14 receiving an encapsulation mapping; and

15 storing the encapsulation mapping as the encapsulation mapping entry in
16 the encapsulation mapping table.
17

18 **81.** The one or more processor-accessible media as recited in claim 76,
19 wherein at least a portion of the processor-executable instructions comprise at
20 least part of a migrator intermediate driver that is positioned below a protocol
21 stack portion of a network stack.
22
23
24
25

1 **82.** The one or more processor-accessible media as recited in claim 76,
2 wherein the source/destination pair of the located encapsulation mapping entry
3 comprises a source/destination address information pair of the located
4 encapsulation mapping entry.

5
6 **83.** The one or more processor-accessible media as recited in claim 76,
7 wherein the flow of packets for the connection comprises a connection that has
8 been migrated.

9
10 **84.** The one or more processor-accessible media as recited in claim 76,
11 wherein the action of extracting comprises an action of:

12 extracting a source internet protocol (IP) address, a destination IP
13 address, a source transmission control protocol (TCP) port, and a
14 destination TCP port from the located encapsulation mapping entry.

15
16 **85.** The one or more processor-accessible media as recited in claim 76,
17 wherein the action of replacing comprises an action of:

18 replacing the flow identifier of the encapsulated packet with a source
19 transmission control protocol (TCP) port and a destination TCP port from
20 the located encapsulation mapping entry to produce the de-encapsulated
21 packet for the connection.

1 86. The one or more processor-accessible media as recited in claim 76,
2 wherein the action of replacing comprises an action of:

3 replacing an originating internet protocol (IP) address and a
4 destination IP address of the encapsulated packet with a source IP address
5 and a destination IP address, respectively, from the located encapsulation
6 mapping entry to produce the de-encapsulated packet for the connection.
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25